(12) UK Patent Application (19) GB (11) 2 359 222 (13) A

(43) Date of A Publication 15.08.2001

(21) Application No 0000237.8

(22) Date of Filing 07.01.2000

(71) Applicant(s)
3Com Corporation
(Incorporated in USA - Delaware)
5400 Bayfront Plaza, Santa Clara,
California 95952-8145, United States of America

(72) Inventor(s)
Ruben Aszkenasy
David John Law
Simon Peter Valentine
Anthony R Newton

(74) Agent and/or Address for Service
Bowles Horton
Felden House, Dower Mews, High Street,
BERKHAMSTED, Herts, HP4 2BL, United Kingdom

(51) INT CL[7]
H04L 12/56 // H04L 12/44 29/06

(52) UK CL (Edition S )
H4P PPBC PPG

(56) Documents Cited
EP 0963079 A1 JP 040284759 A US 5734824 A

(58) Field of Search
UK CL (Edition R ) H4P PEUL PEUM PEUX PF PPBC PPF PPG
INT CL[7] H04L 12/26 12/44 12/56 29/06
Online Databases: WPI, EPODOC, JAPIO, TXTUS1, TXTUS2, TXTEP1, TXTGB1,TXTWO1, INSPEC

(54) Abstract Title

# Method of exchanging address information using auto-negotiation to assist network topology

(57) In the auto-negotiation process between two network devices according to the IEEE 802.3 standard, the 'Next Page' function, which permits additional information to be exchanged during the link-up process, is used to exchange address information. MAC or IP addresses and a code specifying the nature of the address information are transmitted in the Next Page message code fields. The address information may be retrieved from memory and stored in registers under the control of a state machine which performs the auto-negotiation. The address information is used to deduce network topology without the use of a complex deduction algorithm.

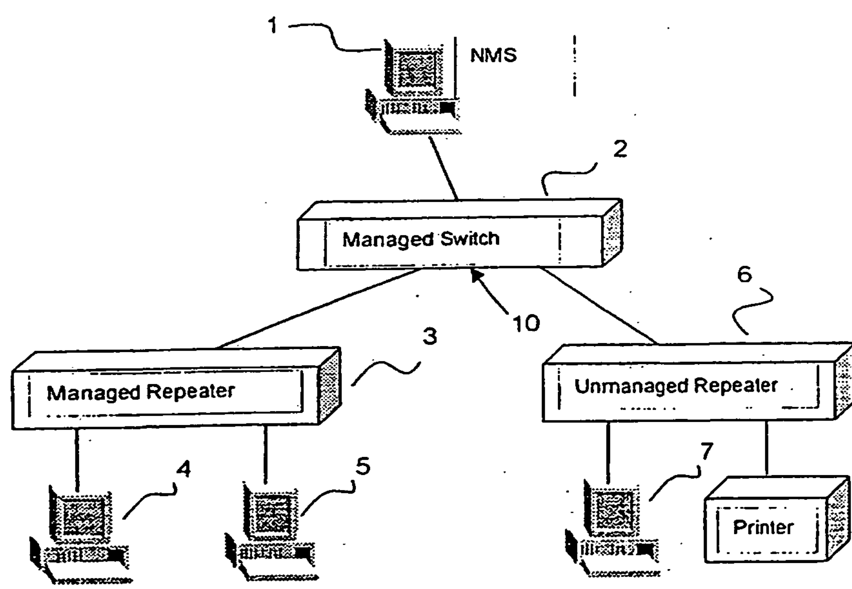


FIG.1

GB 2 359 222 A

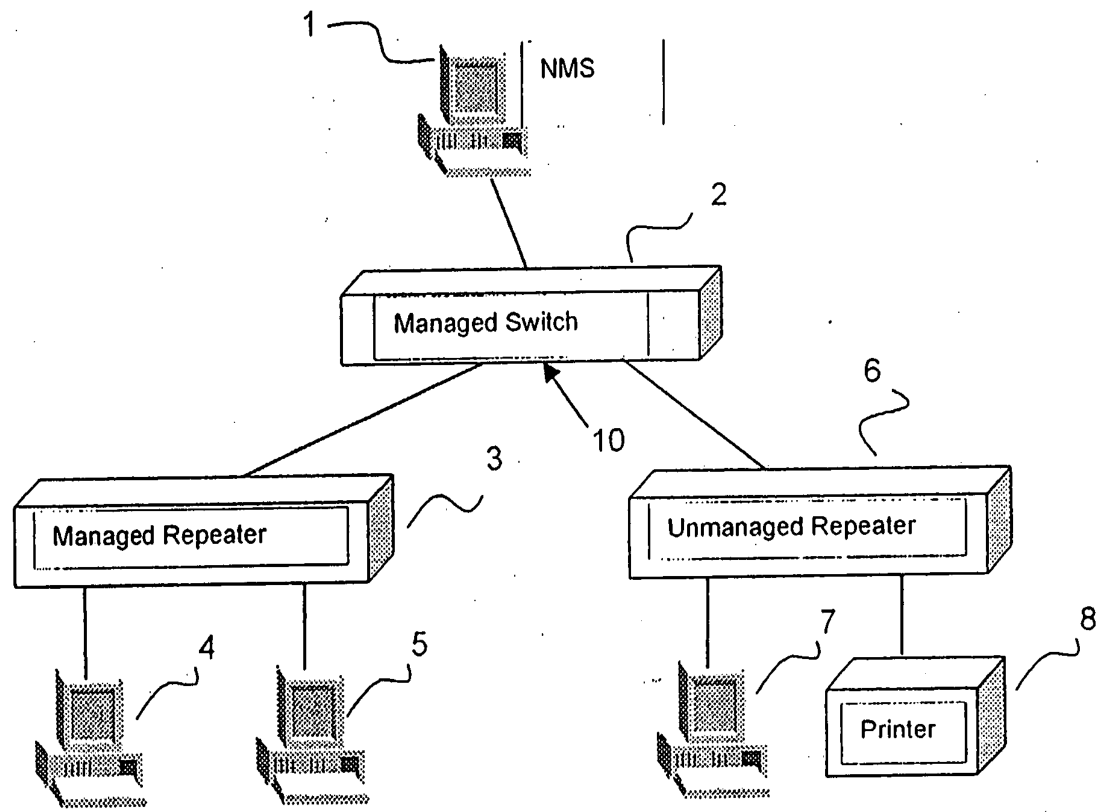
1
NMS
2
Managed Switch
10
3
Managed Repeater
6
Unmanaged Repeater
4
5
7
8
Printer

FIG.1

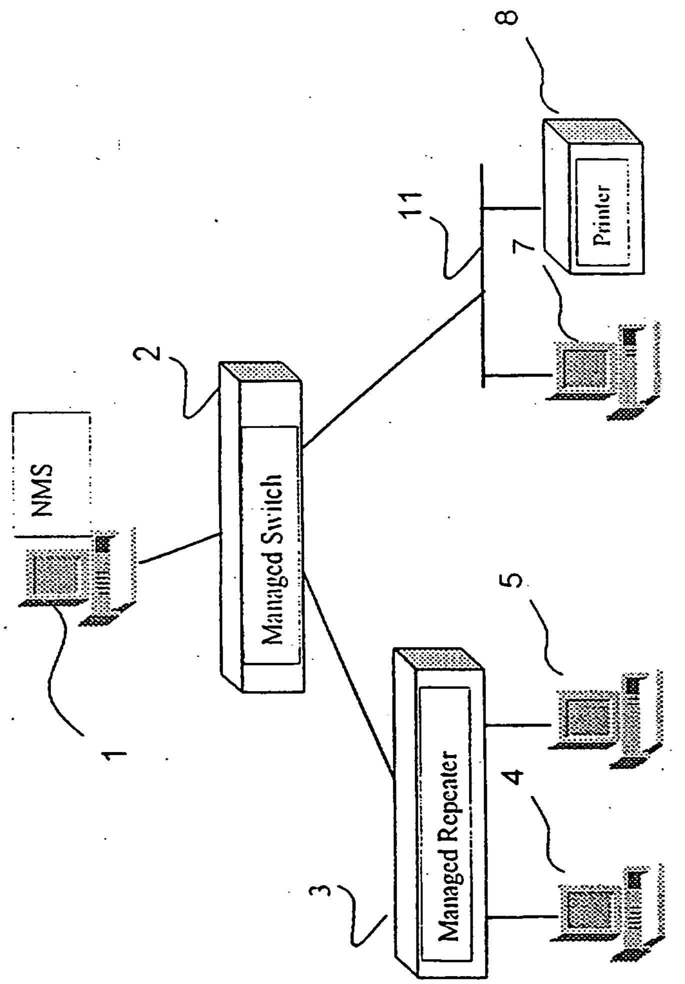


FIG.2

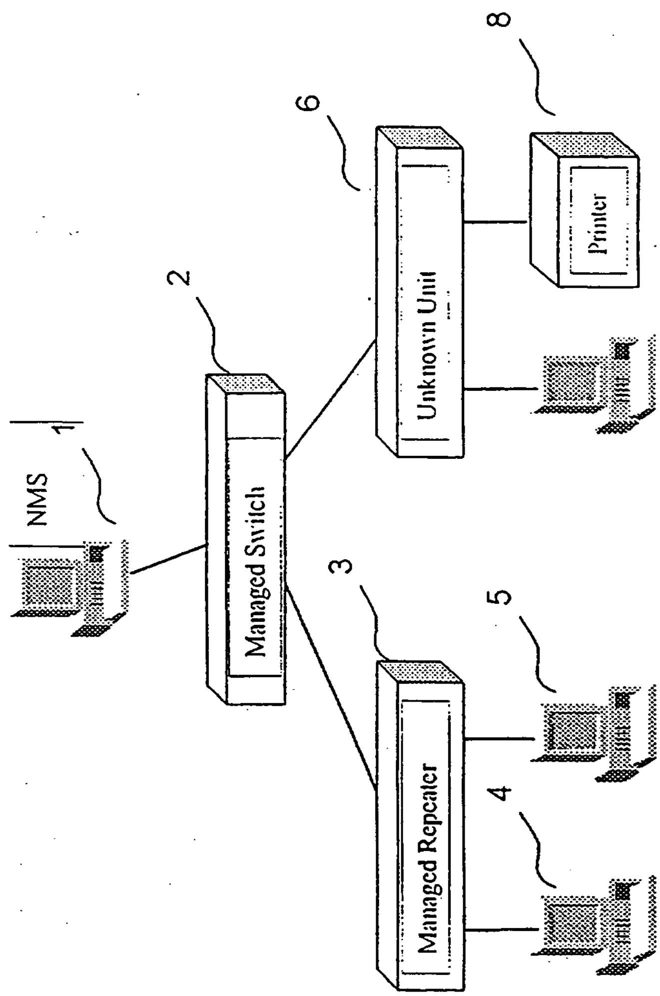

NMS
1
2
Managed Switch
3
Managed Repeater
4
5
6
Unknown Unit
8
Printer


FIG.3

| | 15 | | | 12 | 11 | | | 8 | 7 | | | 4 | 3 | | | 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| message code | NP | Ack | MP | Ack2 | T | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 |
| 1st user code | NP | Ack | MP | Ack2 | T | $O_{23}$ | $O_{22}$ | $O_{21}$ | $O_{20}$ | $O_{19}$ | $O_{18}$ | $O_{17}$ | $O_{16}$ | $O_{15}$ | $O_{14}$ | $O_{13}$ |
| 2nd user code | NP | Ack | MP | Ack2 | T | $O_{12}$ | $O_{11}$ | $O_{10}$ | $O_9$ | $O_8$ | $O_7$ | $O_6$ | $O_5$ | $O_4$ | $O_3$ | $O_2$ |
| 3rd user code | NP | Ack | MP | Ack2 | T | $O_1$ | $O_0$ | $U_{19}$ | $U_{18}$ | $U_{17}$ | $U_{16}$ | $U_{15}$ | $U_{14}$ | $U_{13}$ | $U_{12}$ | $U_{11}$ |
| 4th user code | NP | Ack | MP | Ack2 | T | $U_{10}$ | $U_9$ | $U_8$ | $U_7$ | $U_6$ | $U_5$ | $U_4$ | $U_3$ | $U_2$ | $U_1$ | $U_0$ |

**FIG.4**

| | 15 | 14 | 13 | 12 | 11 | 10 | 9 | 8 | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| message code | NP | Ack | MP | Ack2 | T | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 1 |
| 1st user code | NP | Ack | MP | Ack2 | T | $3ComOUI_{23-13}$ | | | | | | | | | | |
| 2nd user code | NP | Ack | MP | Ack2 | T | $3ComOUI_{12-2}$ | | | | | | | | | | |
| 3rd user code | NP | Ack | MP | Ack2 | T | $3ComOUI_{1-0}$ | | $OP_3$ | $OP_2$ | $OP_1$ | $OP_0$ | $D_{15}$ | $D_{14}$ | $D_{13}$ | $D_{12}$ | $D_{11}$ |
| 4th user code | NP | Ack | MP | Ack2 | T | $D_{10}$ | $D_9$ | $D_8$ | $D_7$ | $D_6$ | $D_5$ | $D_4$ | $D_3$ | $D_2$ | $D_1$ | $D_0$ |

FIG.5

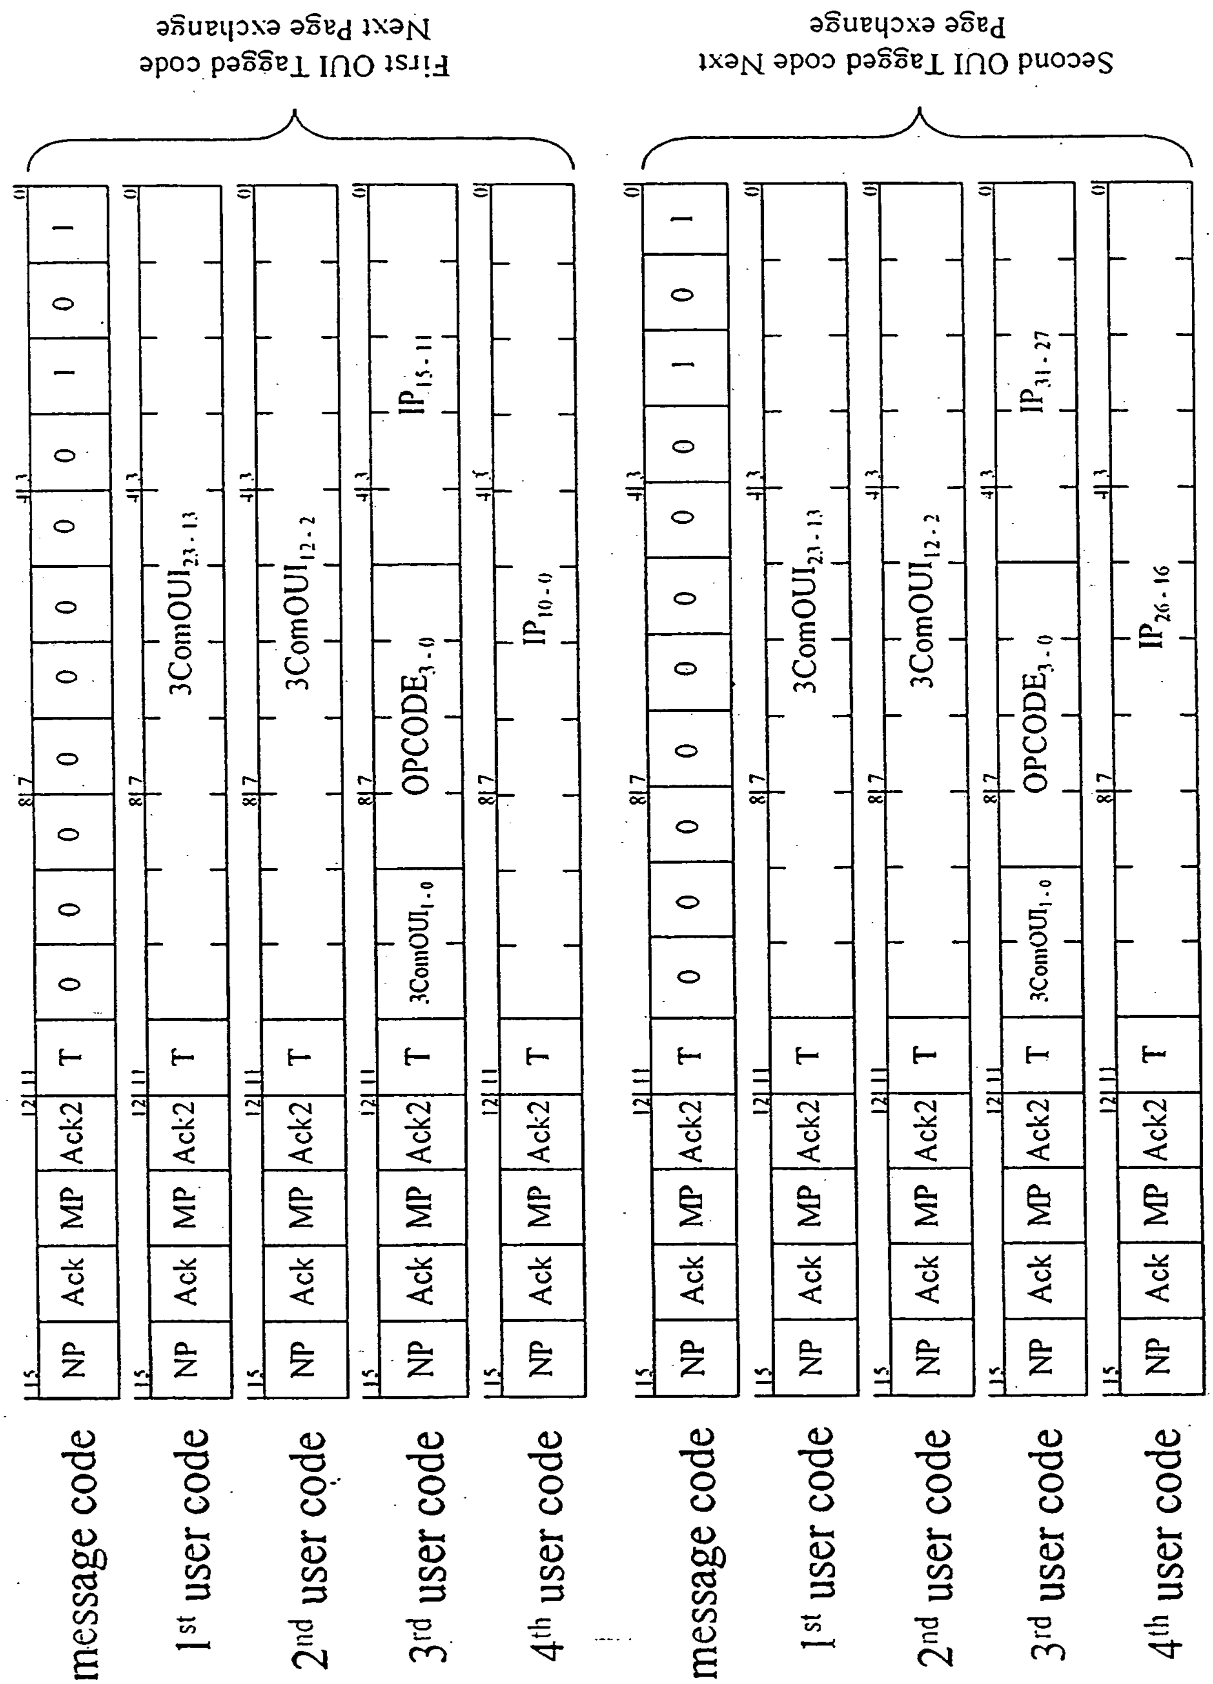
First OUI Tagged code Next Page exchange
Second OUI Tagged code Next Page exchange
message code
1st user code
2nd user code
3rd user code
4th user code
NP
Ack
MP
Ack2
T
3ComOUI23-13
3ComOUI12-2
3ComOUI1-0
OPCODE3-0
IP15-11
IP10-0
IP31-27
IP26-16
15
12 11
8 7
4 3
0

FIG. 6

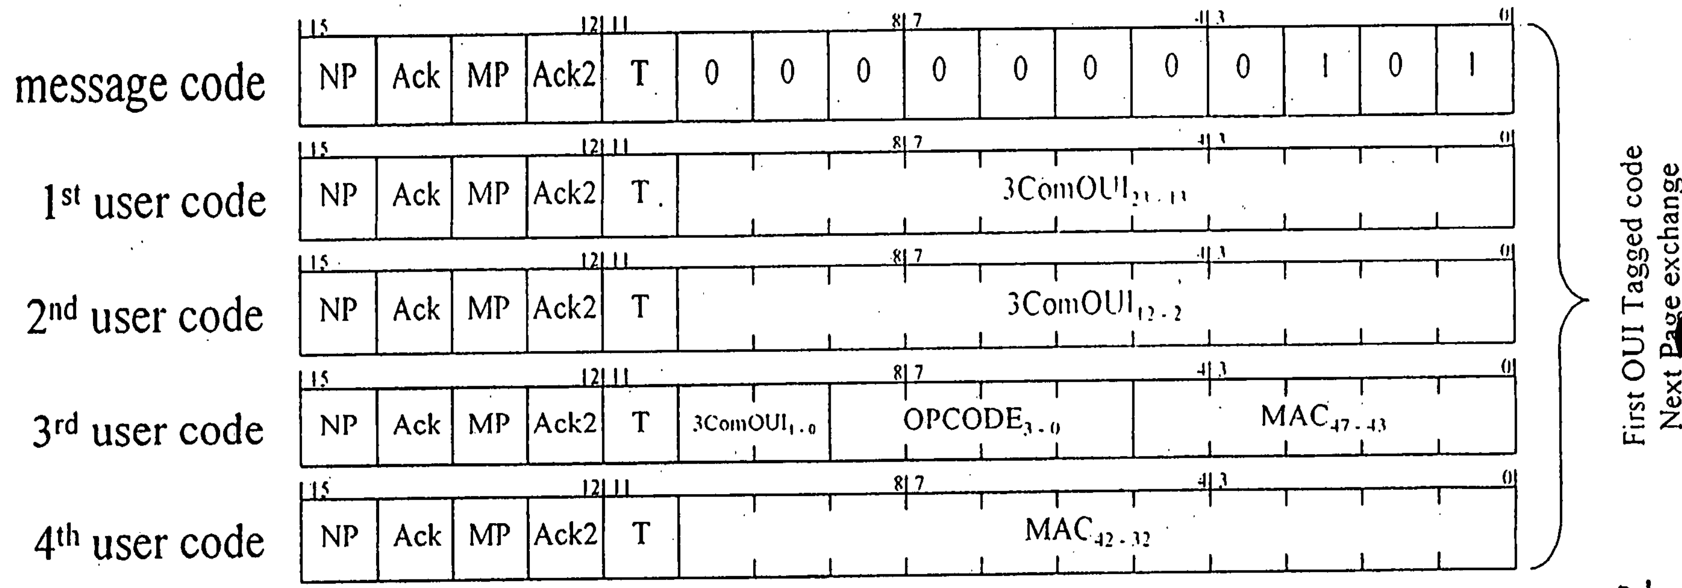

message code
1st user code
2nd user code
3rd user code
4th user code
15
12
11
8
7
4
3
0
NP
Ack
MP
Ack2
T
0 0 0 0 0 0 0 0 1 0 1
3ComOUI21-13
3ComOUI12-2
3ComOUI1-0
OPCODE3-0
MAC47-43
MAC42-32
First OUI Tagged code Next Page exchange


FIG.7A

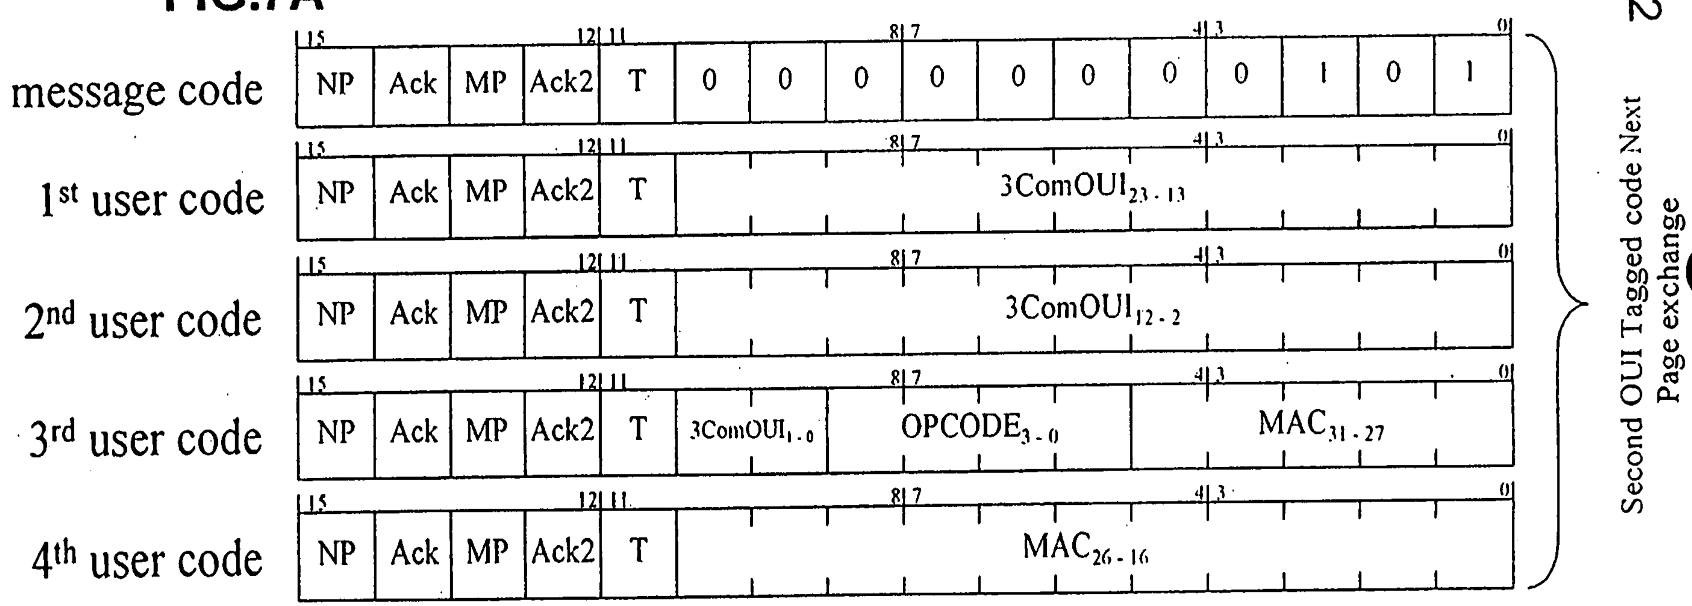

message code
1st user code
2nd user code
3rd user code
4th user code
15
12
11
8
7
4
3
0
NP
Ack
MP
Ack2
T
0 0 0 0 0 0 0 0 1 0 1
3ComOUI23-13
3ComOUI12-2
3ComOUI1-0
OPCODE3-0
MAC31-27
MAC26-16
Second OUI Tagged code Next Page exchange


FIG.7B

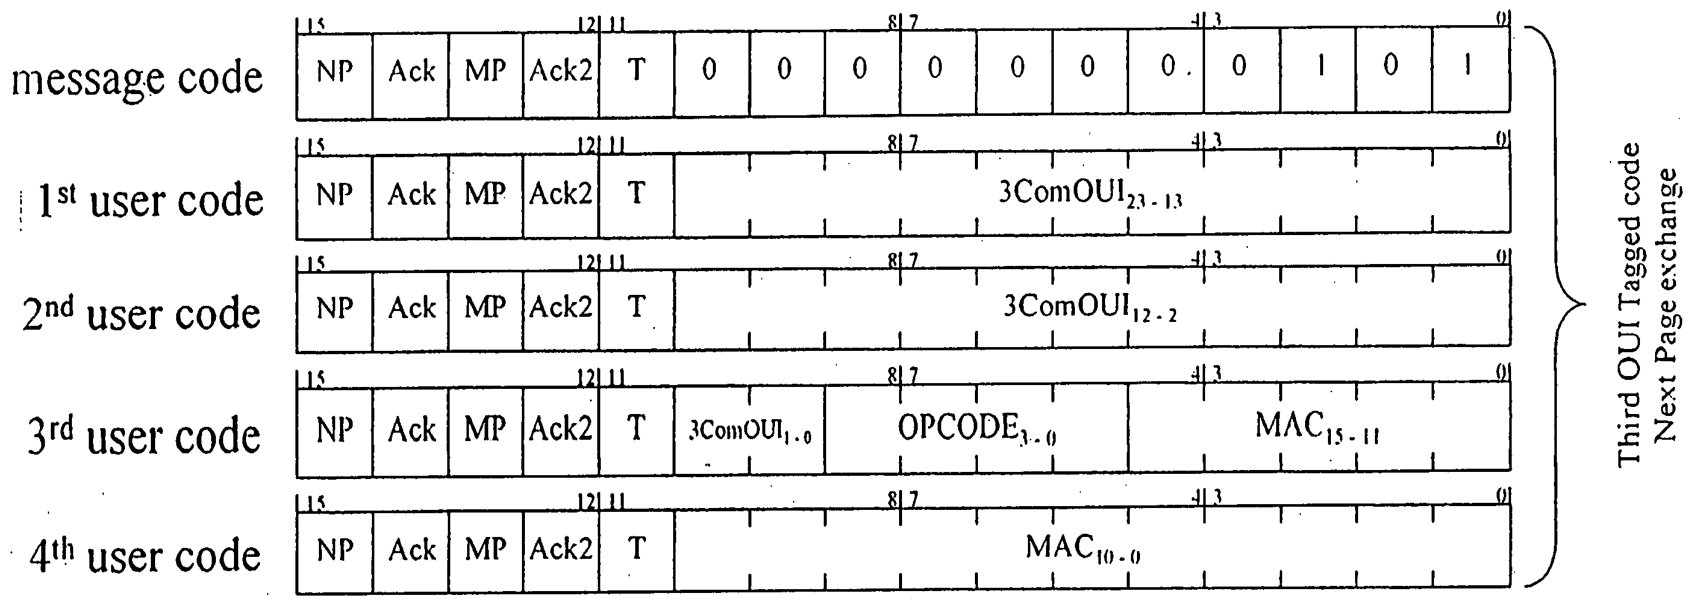
message code
1st user code
2nd user code
3rd user code
4th user code
15
12
11
8
7
4
3
0
NP
Ack
MP
Ack2
T
0 0 0 0 0 0 0 0 1 0 1
3ComOUI23-13
3ComOUI12-2
3ComOUI1-0
OPCODE3-0
MAC15-11
MAC10-0
Third OUI Tagged code
Next Page exchange

FIG.7C

| Op Code | Meaning |
|---|---|
| 0000 | Reserved for future use |
| 0001 | IP Address, Least significant 16 Bits |
| 0010 | IP Address, Most significant 16 Bits |
| 0011 | MAC Address, Least significant 16 Bits |
| 0100 | MAC Address, Mid significant 16 Bits |
| 0101 | MAC Address, Most significant 16 Bits |
| 0110 to 1111 | Reserved for future use |

FIG.8

| | 15 | | | 12 | 11 | | | 8 | 7 | | | 4 | 3 | | | 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| message code | NP | Ack | MP | Ack2 | T | $M_{10}$ | $M_9$ | $M_8$ | $M_7$ | $M_6$ | $M_5$ | $M_4$ | $M_3$ | $M_2$ | $M_1$ | $M_0$ |
| 1st user code | NP | Ack | MP | Ack2 | T | $I_{31}$ | $I_{30}$ | $I_{29}$ | $I_{28}$ | $I_{27}$ | $I_{26}$ | $I_{25}$ | $I_{24}$ | $I_{23}$ | $I_{22}$ | $I_{21}$ |
| 2nd user code | NP | Ack | MP | Ack2 | T | $I_{20}$ | $I_{19}$ | $I_{18}$ | $I_{17}$ | $I_{16}$ | $I_{15}$ | $I_{14}$ | $I_{13}$ | $I_{12}$ | $I_{11}$ | $I_{10}$ |
| 3rd user code | NP | Ack | MP | Ack2 | T | $I_9$ | $I_8$ | $I_7$ | $I_6$ | $I_5$ | $I_4$ | $I_3$ | $I_2$ | $I_1$ | $I_0$ | R |

**FIG.9**

| | 15 | | | 12 | 11 | | | 8 | 7 | | | 4 | 3 | | | 0 |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| message code | NP | Ack | MP | Ack2 | T | $M_{10}$ | $M_9$ | $M_8$ | $M_7$ | $M_6$ | $M_5$ | $M_4$ | $M_3$ | $M_2$ | $M_1$ | $M_0$ |
| 1st user code | NP | Ack | MP | Ack2 | T | $A_{47}$ | $A_{46}$ | $A_{45}$ | $A_{44}$ | $A_{43}$ | $A_{42}$ | $A_{41}$ | $A_{40}$ | $A_{39}$ | $A_{38}$ | $A_{37}$ |
| 2nd user code | NP | Ack | MP | Ack2 | T | $A_{36}$ | $A_{35}$ | $A_{34}$ | $A_{33}$ | $A_{32}$ | $A_{31}$ | $A_{30}$ | $A_{29}$ | $A_{28}$ | $A_{27}$ | $A_{26}$ |
| 3rd user code | NP | Ack | MP | Ack2 | T | $A_{25}$ | $A_{24}$ | $A_{23}$ | $A_{22}$ | $A_{21}$ | $A_{20}$ | $A_{19}$ | $A_{18}$ | $A_{17}$ | $A_{16}$ | $A_{15}$ |
| 4th user code | NP | Ack | MP | Ack2 | T | $A_{14}$ | $A_{13}$ | $A_{12}$ | $A_{11}$ | $A_{10}$ | $A_9$ | $A_8$ | $A_7$ | $A_6$ | $A_5$ | $A_4$ |
| 5th user code | NP | Ack | MP | Ack2 | T | $A_3$ | $A_2$ | $A_1$ | $A_0$ | $R_6$ | $R_5$ | $R_4$ | $R_3$ | $R_2$ | $R_1$ | $R_0$ |

FIG.10

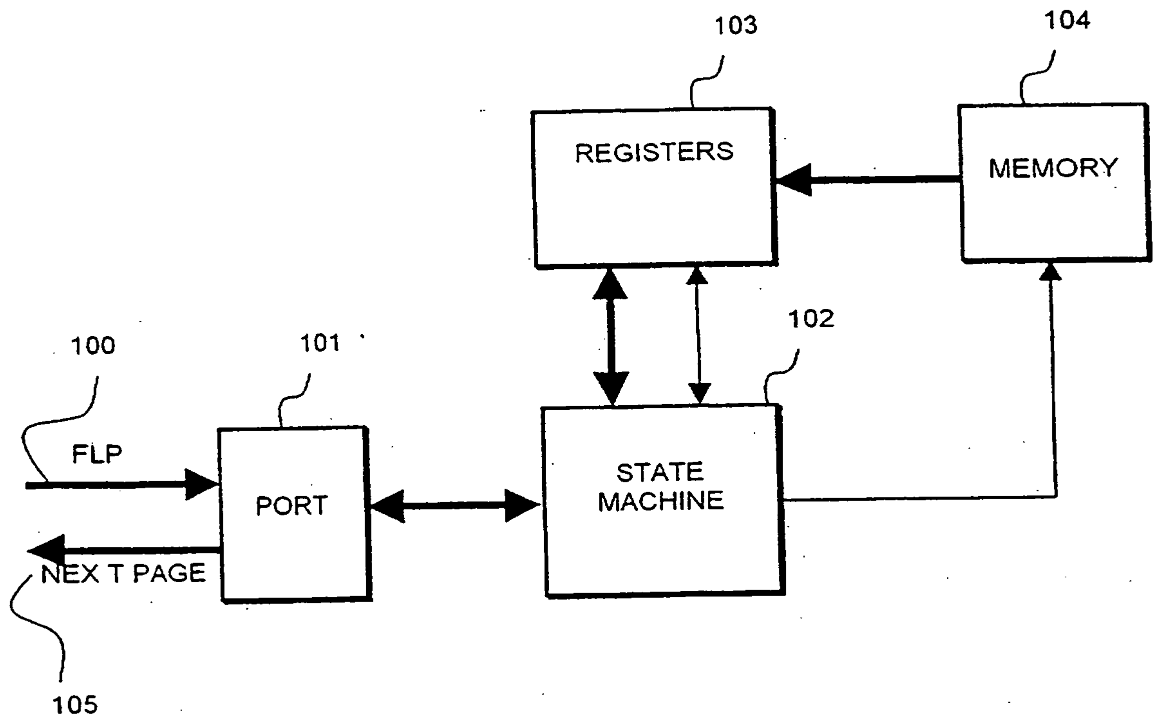
103
104
REGISTERS
MEMORY
102
100
101
FLP
STATE
MACHINE
PORT
NEX T PAGE
105

FIG.11

2359222

# METHOD OF EXCHANGING ADDRESS INFORMATION USING AUTO-NEGOTIATION TO ASSIST NETWORK TOPOLOGY

## Field of the Invention

This invention generally relates to the design and/or management of packet-based communication networks and more particularly to the process of discovery of a network's topology. It specifically concerns the development of a technique known as auto-negotiation to enable the exchange of address information and thereby to facilitate the process of discovery.

## Background to the Invention

Having the ability to represent a computer network makes debugging a network significantly easier. A network administrator would be able to see the network topology and can use that knowledge to aid in performance diagnostics or a re-design of the said network. Knowing the network topology also makes life easier for the network administrator to configure the network, because the adminstrator can traverse the network devices (using say 'Telnet' or web browsers) and can make configuration changes and know what devices those changes may or may not effect.

The present invention is intended to facilitate the process of topology discovery, and particularly to facilitate the exchange of address information (such as network addresses and/or media access control addresses). It also has the major advantage that the IP or MAC address information is exchanged before a link between devices is fully established for the passage of data packets, so that valuable network time or bandwidth is not unnecessarily occupied

During the development of the 100BASE-T standard (802.3u) Auto-Negotiation, the technique of passing configuration information from one device to another as part of the link start-up sequence was defined. This technique provides a means whereby devices at the ends of a 'twisted-pair' link could exchange their abilities and then start-up running at

their highest common operating ability. As part of this auto-negotiation function the additional, optional, 'Next Page' function was defined (see IEEE Standard 802.3, 1998, Clause 28.2.3.4). The 'Next Page' function allows the transfer of arbitrary data between two devices on a link after the basic configuration information has been exchanged but prior to the link going into operation (i.e. the normal transmission of addressed data packets). Subsequently, during the development of the Gigabit Ethernet Standard (802.3z), the Auto-Negotiation function was extended to operate over 1Gb/s fibre links.

The invention is based on the fact that this Next Page transfer can only occur on the point to point link between two PHYs (physical layer devices), and the realisation that this fact enables the unique identification of the device at a far end of the link. Thus the 'Next Page' functional tool forms the basis of a faster and more reliable topology discovery method.

There are (as described in the Standard) two types of Auto-Negotiation Next Page encoding, "Message" Pages and "Unformatted" or "User" Pages. A Next Page message exchange consists of the exchange of a Message page and a number of User/Unformatted pages. The Message page defines the type of Next Page exchange taking place by the message code it contains. The number of User/Unformatted pages that follow is determined by the particular Next Page message code.

There are at least two approaches that can be used to provide the transfer of IP and/or MAC addressees between the two PHYs at the end of a link. The first is to use the ability to transfer proprietary user defined information as part of the already defined Message code #5 - Organisationally Unique Identifier (OUI) tag code, Next Page exchange (see IEEE 802.3-1998 Clause 28C.6). Another approach is to define new Next Page tag codes that provide for exchange of MAC/IP addresses directly. The second approach would be preferable, because since it could provide a multi-vendor interoperable system that uses fewer Next Pages exchanges than the OUI tag code approach. However currently it would have the disadvantage that it would require allocation of codes by the authority (IEEE) controlling the Standard.

Brief Description of the Drawings

Figure 1 illustrates a simple packet-based communication network.

Figure 2 illustrates the network as it is interpreted by a network management station.

Figure 3 illustrates the network as it may be interpreted employing the invention.

Figures 4 to 10 illustrate various forms of 'Next Page' messages which can be used to perform the invention.

Figure 11 illustrates schematic the relevant parts of a device which incorporates the invention.

Description of a Preferred Example

Figure 1 shows a simple network wherein an NMS (network management station) 1 is connected to a managed switch 2 connected by one port to a managed repeated 3 itself connected to 'users' (personal computers) 4 and 5. By another port the managed switch 2 is connected to an unmanaged repeater 6 itself connected to a PC 7 and a printer 8.

The responsibility of the NMS is to monitor and control agents. An agent is a software component residing in a network device e.g. a switch or a repeater. Devices that can run agent software are known as 'managed'. Conversely, unmanaged devices do not run any agent software and are not at present visible to the NMS.

By reading the source media access control address (MAC) addresses of incoming packets, a managed device can learn which MAC addresses are related to a particular port. In Figure 1, the left-hand port 10 of the managed switch 2 will have three MAC addresses associated with it: the managed repeater 3 and its two PCs 4 and 5. The port/MAC data are in accordance with well-known practice stored in agent look-up tables.

By reading the port/MAC look-up tables in each managed device, the NMS attempts to deduce the topology of the network. Figure 2 below shows the network topology from the point of view of the NMS.

Since the unmanaged repeater 6 does not have its own MAC device it is invisible to the NMS, and only the end MAC devices (the PC and the printer) can be detected by the NMS. As far as the NMS is concerned, the right hand connection to the managed switch 2 looks like a wire segment 11.

Since a topology discovery algorithm is carried out by the NMS 1, only it has a view of the network. Although the agents in the network appliances store port/MAC tables, they cannot use this information to deduce their local connectivity. Thus in Figure 2, the managed switch 2 is not aware of the managed repeater 3.

As agents are not aware of their local environment, a fairly complex deduction algorithm is needed by the NMS to discover the overall topology.

Devices learn MAC addresses, e.g. by storing in a forwarding database the source address (SA) of an incoming packet and the allotted number of the port by which the packet was received. However, MAC addresses learnt in this manner can be unreliable. When for example a PC or a printer is disconnected, it takes time for the agent to realise that its port/MAC tables are no longer valid. Moreover, some older types of device do not erase old MAC addresses that they have learnt. This makes it difficult for the NMS to track the movement of, for example, PC's from one port of the network to another.

Furthermore a discovery process generates a large amount of network traffic and therefore may consume excessive time or bandwidth.

When two devices (e.g. a PC and a repeater) first attempt to connect to each other, their physical layer components (PHYs) use a process called 'auto-negotiation'. This is fully described in IEEE Standard 802.3, 1998 Edition, published by the Institute of Electrical and Electronics Engineers, Inc., New York, USA (ISBN 0-7381-0330-6), pages 689 to

734. The main purpose of auto-negotiation is to reconcile the wire data speed between the two devices. Auto-negotiation also has a feature called 'Next Page' that permits additional data to be exchanged during the linking-up process. If present, this extra data is stored in certain PHY registers (as described in the Standard) and can be made available to an agent.

The present invention specifically concerns on using the Next Page function to exchange a unique port identification (ID) that fully and reliably identifies a device and its location. The port ID preferably comprises an address, i.e. a MAC or IP address, a 'Unit number', which may be useful for locating a unit in a large stack in a wiring closet, and a 'hardware port number', which may be useful for identifying a cable connection on a unit.

The agent in a device can store this information for each attached device in a look-up table. Since the address information is directly associated with the PHY connection at each end of the cable, topology deduction is not required, and the agent is now aware of its nearest neighbours. It is now much easier for the NMS to record the exact network topology and to track PCs etc which are moving in the network. The 'learning' problem disappears.

In the event that the nearest neighbour of a managed device is an unmanaged device, the nearest neighbour's 'Next Page' will be empty. The PHY in the managed appliance will register this fact, and therefore the agent in the managed device will still know that there is a unit connected between it and further downstream devices. Thus as shown in Figure 3 the software agent in the managed switch 2 will now be aware of the existence of an unknown unit (actually the unmanaged repeater) 6 between the switch 2 and the devices 7 and 8.

Since the information exchange is carried out during the physical layer link-up process (i.e. not using LAN data packets) valuable network bandwidth is not wasted.

There is an option whether to exchange MAC addresses or protocol (network) addresses (IP addresses) using the Next Page feature.

Whereas MAC addresses are more 'reliable' because not all end devices have IP stacks (e.g. printers), and it takes time for PCs to render their IP stacks fully operational, one cannot use Web agents to 'hot-link' from device to device using a web browser in the NMS. Users would need to use RARP (reverse address resolution protocol) to obtain IP addresses.

Using a Web browser one can hot-link from device to device (with embedded Web agents), but that requires IP addresses which are not necessarily always possessed by some devices.

As can be seen above the exchange of either a IP or MAC address through Next Page exchange can be used for topology determination. The following describes preferred mechanisms to exchange either of these types of addresses. If desired only one of the several possible types of exchange need to be implemented.

The 48-Bit Universal LAN MAC Address, used in Ethernet as the Ethernet MAC Address, consist of two parts (see IEEE802-1990 subclause 5.2). The first 24 bits correspond to the Organisationally Unique Identifier (OUI) as assigned by the IEEE. The second part, comprising the remaining 24, is administered locally by the assignee. When a device is manufactured the OUI is combined with a 24 bit, locally administered value which is registered as used so it can never be used again. In this way a unique 48 bit MAC address is formed. As can be seen however once 2^24 (nearly 17 million) similar units are manufactured the assignee has to return to the IEEE and request a new OUI.

The OUI Tagged Next Page exchange allows the transfer of an OUI, and related user-defined user codes, from a far end device. The transfer of these user-defined user codes within the OUI Tagged Next Page exchange can be used to convey the actual MAC or IP Address information between the two PHYs on the point to point link. As the user-defined user codes are related to the OUI, that is, only if the OUI is known to the receiving device, can the user-defined codes transferred be interpreted, this mechanism is not a multi-vendor approach. If for example another vendor's equipment were to receive an OUI Tagged

Next Page message from a particular vendor, even if it was able to recognise the OUI as belonging to that particular vendor, it would be unable to decode it, because format is defined by that particular vendor.

A related issue is the OUI to use. If the OUI of the equipment in which the PHY is installed were used this would present a challenge whenever a new OUI had to be obtained from the IEEE as described above. As there is no way of predicting what this new OUI may be this would mean that equipment manufactured prior to the new OUI being allocated would be unable to recognise the new OUI and this mechanism would fail. There is however no requirement that the OUI sent in the OUI Next Page message exchange be the same as the OUI of the equipment that the PHY is installed in, only that it is an OUI owned by the manufacturer of the equipment. The OUI therefore used in the OUI Tagged Next Page exchange is an OUI that is already in use, and will never be changed once allocated. This does present an interesting feature for MAC address transfers as it therefore requires that two OUIs are transferred, the OUI Next Page exchange OUI and the OUI of the equipment. This means that a full 48 bit Ethernet MAC address has to be transferred as well as the OUI in each of the OUI Next Page messages.

There is an optimisation when the OUI of the OUI Tagged Next Page message matches the OUI of the equipment sending the massage allowing only 24 bits to be sent but that optimisation is not examined in detail here.

Each message code field employed in the current Standard (see pages 699-702 thereof) is a two-byte message including an 11-bit field and a 5-bit 'header'. Each exchange according to the invention comprises a first field (identified as the 'message code' in each of the remaining figures followed by four subsequent fields, the first to fourth 'user codes'. The message field indicates by its coding (bits 0 and 2) that the following fields are an exchange of information according to the invention and the user codes are as described in the following. The 'header' of each field has five bits (NP, Ack, MP, Ack2 and T) which have the significance prescribed in the Standard. NP indicates whether the Next Page is the last page (0) or not (1). Ack when set to "1" indicates that the device has successfully received its link partner's link code word (see Clause 28.2.1.2.4 of the

Standard). Ack2 when set to "1" indicates an ability to comply with a Next Page message. MP denotes a Message Page (1) or an unformatted page (0). Toggle (T) is a bit that should take the opposite value of a toggle bit in a previously exchange Link Code Word.

Currently, the Standard prescribes two kinds of 'message code pages', each of which consists of a message code and four user codes. The two types are a 'remote fault' page and an OUI Tagged Next Page. The form of the invention shown in Figures 4 to 7 uses an OUI Tagged Next Page.

In Figure 4, bits $0_{23}$-$0_0$ constitute the organisationally unique identifier (OUI). Bits $U_{19}$-$U_0$ are the specific user defined code value. In Figure 5, the bits 3Com $OUI_{23\text{-}0}$ are an assigned 24-bit OUI. Bits $OP_{3\text{-}0}$ are operation code and $D_{15\text{-}0}$ are a 16-bit user defined code value. In Figure 6, $OPCODE_{3\text{-}0}$ is an operation code and $IP_{32\text{-}0}$ is a 32-bit IP address. In Figures 7a to 7c, $MAC_{47\text{-}0}$ is a 48-bit MAC address. In Figure 9, $I_{31\text{-}0}$ is a 32-bit Internet Protocol Address. R represents a bit reserved for future use; it will be set to '0' and ignored on receive. In Figure 10, $A_{47\text{-}0}$ is a 48-bit MAC address. $M_{10\text{-}0}$ indicates a message code (which would require an actual value to be allocated by the authority controlling the Standard).

As will be apparent from, for example, Figure 4, after the inclusion of the OUI bits, $0_{23}$ – $0_0$, only twenty bits ($U_{19}$-$U_1$) remain, so that the maximum 'payload' is only twenty bits per Next Page. It is preferable to allot some of these bits to an operation code (e.g. bits $OP_3$ to $OP_1$ in Figure 5) which can among other things define the nature of the code bits that follow; it may denote whether the bits are part of an IP address, as in Figure 6 or a MAC address as in Figures 7a-7c.

After the allocation of message space (in these examples four bits) to the operation code, 16 bits remain. Since IP addresses have 32 bits and MAC addresses 48 bits, these will need respectively two and three Next Pages for transmission. Thus Figure 6 shows two Next Page exchanges and Figures 7a to 7c show three such exchanges.

Now, as described above, the transfer of the user-defined user codes within the OUI Tagged Next Page exchange is used to convey the actual MAC or IP Address information between the two PHYs on the point to point link.

As multiple OUI Tagged Next Page exchanges must occur to transfer the MAC/IP address information, and since the OUI Next Page message may also be used for other types of proprietary information exchange, further bits of the OUI Next Page message are reserved for identifying the type of proprietary exchange that is taking place. Figure 5 illustrates the assignment of the 20 bits of user defined information. The four most significant bits ($OP_3$ to $OP_0$) are the operation code; the remaining 16 bits provide information specific to the Op-Code. Figure 8 illustrates the different Op-Code assignments to enable the exchange of IP or MAC addresses.

There are two optimisations possible here. The first is to use a different assigned OUI per message type, an OUI for MAC address transfer and a different OUI for IP address transfer. This would enable a reduction in the number of Op-Code bits required. The second optimisation is that once the 16 possible Op-codes are used up a different OUI Tagged Next Page exchange OUI can be used to expand the possible Op-Codes further. While these two optimisations are possible for corporations that have several OUIs allocated to them over time this may not be an option for smaller companies as they may only ever have one OUI allocated to them.

The actual exchange would occur as follows. When the Auto-Negotiation Base Page exchange is complete the Next Page exchange, if supported would commence. Assuming that the far end supported this feature a Next Page OUI Tagged exchange would commence and the four Next Pages would be transferred. Once it is identified that an OUI Tagged Next Page exchange has occurred the OUI supplied would have to be examined. If the OUI were a known OUI, then the Op-Code field would be examined. If the Op-Code indicated an IP/MAC address transfer then the IP/MAC address information can be extracted. This would provide part of the MAC/IP address. A second, and in the case of a MAC address a third, OUI Tagged Next Page exchange occurs, qualified as above, to provide the remainder of the MAC/IP address.

In the case of an IP address which is 32 bits long, the first 16 bits will be supplied in one OUI Tagged Next Page exchange, the remaining 16 bits in a second OUI Tagged Next Page exchange (see Figure 6). In the case of a MAC address which is 48 bits long, the first 16 bits will be supplied in one OUI Tagged Next Page exchange, the next 16 bits in a second transfer and the remaining 16 bits in a third OUI Tagged Next Page exchange (see Figures 7a to 7c). In all cases the type of data transfer, IP or MAC address, is indicated by the Op-Code.

If Auto-Negotiation should complete without the exchange of the part of the MAC/IP Address being supplied, say for example only one of the OUI Tagged Next Page exchanges takes place, then an error should be posted and the MAC/IP Address cannot be considered received.

The second approach is to define a new Next Page Message Code to transfer the address information directly. Both an IP and a MAC address tag code could be defined (the actual addition of a Next Page Message Code to the IEEE 802.3 Standard is controlled by the IEEE so the actual assignment of the new Next Page Message Code would have to pass through the normal IEEE process).

The Internet Protocol (IP) tag code Next Page exchange (see Figure 9) would consist of four Next Page exchanges. The first next page would consist of the Message page containing the IP Tag Message Code ($M_{10}$ to $M_0$) and, the following three messages would contain the actual 32 bit IP Address ($I_{31}$ to $I_0$).

The MAC Address (MAC) tag code Next Page exchange (see Figure 10) would consist of 6 Next Page exchanges. The first next page would consist of the Message page containing the MAC Tag Message Code, the following five messages would contain the actual 48 bit MAC Address ($A_{47}$ to $A_0$).

The payload of these new Next Page messages would only contain the IP or MAC Address information. Hence as the entire payload would be defined in a Standard any

vendor would be able to read the IP/MAC address contained in these messages regardless of the vendor of the equipment at the far end of the link. This would allow multi-vendor interoperable implementations of this topology discovery. This is not possible with the current OUI Tagged Next Page message Code as the OUI in the message has to be known before the used defined information can be read.

In addition, in both cases these new types of Next Page exchange also provide an improvement in performance. Using the OUI Next Page exchange to do a MAC or IP address exchange takes two or three OUI Tagged Next Page exchanges respectively, a total transfer of ten Next Pages in the case of an IP address and fifteen Next Pages in the case of a MAC address. In the case of the new IP address Next Page exchange this is reduced to four pages; in the case of MAC address Next Page exchange, six pages.

Figure 11 shows for the sake of completeness in simplified form part of a network device incorporation the invention. Fast Link Pulses (FLP) are received on an inward path 100 by a port 101 which was an associated physical layer device (PHY) including a state machine 102 organised to perform auto-negotiation as described previously. In normal operation of the state machine values held in registers 103 will be transferred to a remote device by auto-negotiation. The format and timing of the Next Pages exchanges 105 is determined by Fast Link Pulses. In the present invention, the relevant identification information, e.g. the OUI, MAC address and IP address as the case may be is transferred under software control (which also may govern the state machine) to relevant registers within registers 103 and incorporated as previously described in Next Page exchanges controlled by the state machine.

Claims

1. A method of exchanging identifying information between two devices which are capable of auto-negotiation according to IEEE Standard 802.3, comprising transmitting a plurality of Next Page message code fields which include fields constituting address information and a code specifying the nature of the address information.

2. A method according to claim 1 wherein the address information comprises a media access control (MAC) address.

3. A method according to claim 1 wherein the address information comprises a network (IP) address.

4. A method according to any foregoing claim and further comprising retrieving said address information from memory, and storing said address information in registers under the control of a state machine which performs said auto-negotiation.

5. A device for use in a packet-based communication system and including at least one port associated with a state machine for performing auto-negotiation including Next Page exchanges, said Next Page exchanges including fields constituting an address of said device and at least one code field indicating the nature of said address, and means for providing said address for incorporation in said Next Page exchanges.

| | | | |
|---|---|---|---|
| **Application No:** | GB 0000237.8 | **Examiner:** | Gareth Griffiths |
| **Claims searched:** | All | **Date of search:** | 23 June 2000 |

13

**Patents Act 1977**
**Search Report under Section 17**

**Databases searched:**

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.R): H4P (PEUL, PEUM, PEUX, PF, PPBC, PPF, PPG)

Int Cl (Ed.7): H04L 12/26, 12/44, 12/56, 29/06

Other: Online Databases: WPI, EPODOC, JAPIO, TXTUS1, TXTUS2, TXTEP1, TXTGB1, TXTWO1, INSPEC

**Documents considered to be relevant:**

| Category | Identity of document and relevant passage | | Relevant to claims |
|---|---|---|---|
| A | EP0963079 A | (ADVANCED MICRO DEVICES) | |
| A | US5734824 | (CHOI) | |
| A | JP040284759 A | (NEC SOFTWARE) | |

| | | | |
|---|---|---|---|
| X | Document indicating lack of novelty or inventive step | A | Document indicating technological background and/or state of the art. |
| Y | Document indicating lack of inventive step if combined with one or more other documents of same category. | P | Document published on or after the declared priority date but before the filing date of this invention. |
| | | E | Patent document published on or after, but with priority date earlier than, the filing date of this application. |
| & | Member of the same patent family | | |